

White Paper Report

Report ID: 98553

Application Number: HD5020307

Project Director: Colin Allen (colallen@indiana.edu)

Institution: Indiana University, Bloomington

Reporting Period: 9/1/2007-12/31/2008

Report Due: 3/31/2009

Date Submitted: 9/15/2010

Running Head: Toward dynamic representation of philosophy

Full Title: From encyclopedia to ontology: Toward dynamic representation of the discipline of philosophy

Abstract: The application of digital humanities techniques to philosophy is changing the way scholars approach the discipline. This paper seeks to open a discussion about the difficulties, methods, opportunities, and dangers of creating and utilizing a formal representation of the discipline of philosophy. We review our current project, the Indiana Philosophy Ontology (InPhO) project, which uses a combination of automated methods and expert feedback to create a dynamic computational ontology for the discipline of philosophy. We argue that our distributed, expert-based approach to modeling the discipline carries substantial practical and philosophical benefits over alternatives. We also discuss challenges facing our project (and any other similar project) as well as the future directions for digital philosophy afforded by formal modeling.

DRAFT VERSION—DO NOT CITE

(Comments welcome—send to cbuckner@indiana.edu)

From Encyclopedia to Ontology: Toward Dynamic Representation of the Discipline of Philosophy

1. Introduction

Encyclopedias have always occupied a precarious position in academia. On the one hand they taxonomize human knowledge and provide valuable entry points for scholars and students into the intellectual worlds of academic disciplines, covering their subject matters in more breadth and detail than could any single person or even any reasonably-sized university department. On the other hand, they carry a risk of congealing knowledge into a cold and quickly-obsolete imitation of living scholarship, stultifying the thought of beginners who might be better off wrestling with multiple, recent perspectives than the predigested orthodoxy of a designated expert. The tension in being encyclopedic is especially acute today, given the recent explosion in the number of universities, scholars, and academic publications. While the explosion has made faithful and succinct summarization even more elusive, encyclopedias have perhaps never been more relevant. If we wish to prevent disciplines from disintegrating into collections of highly-technical cottage industries in which specialists speak only amongst themselves, the development and maintenance of reference works offering accessible, up-to-date summaries is imperative.¹

Fortunately, the advent of the digital humanities has brought a rich new arsenal of strategies to help us respond intelligently to the academic avalanche. The development of dynamic reference works such as online encyclopedias has brought some relief, as these

publications are more responsive to new work than was possible with traditional printing methods. Different editorial models have been created for dynamic reference works, from Wikipedia's open authorship and editorship to the Stanford Encyclopedia of Philosophy's more traditional editorial policies of selected editors reviewing the work of invited authors. Even though the availability of online encyclopedias does, in principle, make the latest scholarship more accessible, the relatively unsophisticated capacities of search engines and their users make it likely that the full potential of digital encyclopedias is far from being realized.²

In short, scholars and students don't just need the reference works—they also need the means to search and navigate them effectively. To preserve the utility of encyclopedias as they grow, we must also improve our ability to represent their contents in meaningful ways accessible to novice and expert alike. The dynamic nature and increased scale of digital reference works, however, render traditional editorial methods of gathering and organizing metacontent (indices, cross-references, tables of contents) so resource-intensive and inefficient as to be practically inapplicable. This is especially true for projects with limited staff and resources, common consequences of adherence to the ideals of open access. To address these problems adequately, more sophisticated techniques of generating metacontent from large, asynchronously-updated textual corpora are required. These issues should be addressed with the help of domain experts and not merely by technologists working according to their own ideas of what scholars might need. The best practices require technological expertise to design the formal representations and domain expertise to capture the semantic content of corpora. Decisions here should not be made lightly, as choices in the process of representation embody substantive commitments about the nature of the subject which, if widely-adopted, may come to affect the trajectory of the discipline itself.

2. The Indiana Philosophy Ontology Project

This paper seeks to open a discussion about the difficulties, methods, opportunities, and dangers of creating and utilizing a formal representation of the discipline of philosophy. We review our current project at Indiana University, the Indiana Philosophy Ontology (InPhO)³ project. We are developing an ontology for the Stanford Encyclopedia of Philosophy (SEP),⁴ an online, open access, dynamic reference work, which should be suitable for deployment in other digital philosophy applications.⁵ Our approach to representing the discipline derives from what computer and information scientists call “formal ontology.” However, we will avoid using this term because some researchers sensitive to Husserl’s distinction between “formal ontology” and “material ontology” prefer to reserve it use for a stricter notion of ontology than is satisfied by some aspects of our approach (Poli 1995; Arp this issue; see 6.2 below). Instead, we prefer to refer to our representation as a “computational ontology” and “dynamic ontology.”

We speak of ‘ontology’ in the information science sense of the term, rather than the “metaphysical” sense which may be more familiar to some readers. In information science, ‘computational ontology’ denotes a formally-encoded specification of the concepts relevant to a subject domain (including their properties and relations between them) and a hierarchical classification of those concepts into categories and subcategories (Noy and McGuinness 2001; Gruber 1993, 2008).⁶ The purpose of such an ontology is to assist humans and automatic agents in understanding the contents of the domain (especially in terms of properties, relations, and subsumption/inheritance relationships which hold between the domain’s types) and to allow data generated in one project to be interoperable with others.

2.1 Some Nuts and Bolts: The SEP and Why It Needs an Ontology

The SEP is the first attempt by academics of any discipline to organize their own professional subject matter by collaboratively writing, publishing, and maintaining a dynamically-updated reference work entirely on the web. Articles are submitted by more than 1000 volunteer authors and asynchronously updated to reflect the latest in scholarship. As of January 2009, those authors have submitted almost 1,100 entries (with an additional 300+ currently on commission) containing almost 12.5 million words. Everything submitted by the authors is in turn reviewed by the over 115 volunteer subject editors before publication on the web (a critical difference between the SEP's approach and, for example, wiki-style approaches). As a result, the SEP is both authoritative and comprehensive, as evidenced by its averaging almost 600,000 entry downloads per week. Consistently topping the Google™ search lists for philosophical concepts and thinkers, the SEP has emerged as the most visible and popular online reference work for the discipline of philosophy.

Since its inception, a major goal of the SEP has been to keep the encyclopedia available without charge both to scholars and the public. This goal has so far been satisfied through the volunteer efforts of many field experts, grants from federal and other sources, and a major fundraising effort involving the international community of librarians. The innovative nature of the work, however, brings with it a host of new difficulties not faced by traditional encyclopedias. It is increasingly impractical, for instance, to have editorial staff manually manage cross-references, tables of contents, search keywords, and other metacontent due to the asynchronous submission and revision of articles. There is also a pressure to minimize the editorial burden placed on volunteer contributors, who cannot be expected to constantly monitor the massive, ever-changing contents of the SEP and update metadata themselves. These challenges, coupled with the desire to preserve free, open access to the encyclopedia, create a

strong drive to develop automated and semi-automated information-management tools which can be integrated into the editorial workflow of the encyclopedia. The development of an ontology for the domain of philosophy is central to the success of these tools.

2.2 The InPhO: A Dynamic Ontology

In the design, implementation, and long-term deployment of computational ontologies, knowledge modelers face several enduring challenges. For one, computational ontologies have often been designed without sufficient logical rigor, which may come with pragmatic costs in terms of the expressive power, clarity, and interoperability of the scheme (Guarino 1995; Smith, 2006; Arp, this volume). The economics of ontology design is also a problem, generally requiring significant time from scholars specially-trained in both the target domain and the principles and methods of computational ontology design (hereafter, “double experts”). Obsolescence looms large, as change in the problem domain or our understanding of it can render all that design effort useless, in the best case requiring more time from double experts to manually evolve the ontology (Flouris et al 2006; Ceusters & Smith 2006) and in the worst cases taking a project back to the drawing board.

One broad response to these challenges, emphasized by the “formal ontologists,” is to attempt to produce a “once and for all” description of the underlying reality of the subject domains, and to link the types of those subject domains into a standardized upper-level ontology describing the most basic, enduring features of reality. While this approach can hope to minimize the amount of change needed in future iterations, when change is called for it is usually performed manually. Another approach might be characterized by the phrase “dynamic ontology.” On this approach, more effort is placed on automating as much of the design and

evolution process as possible rather than on attempting to produce a final description in the initial stages of a project.

Dynamic ontology has a slightly different problem space than other approaches. For one, forgoing the use of double experts is a double-edged sword; while dynamic ontology can aspire to be more economical than alternatives, it must be more creative in its methods of obtaining data in ontology construction. The automatic processing of heterogeneous sources of data (often of different degrees of reliability) will often be required, and the problems of data inconsistency and validation loom large. Automated methods of ontology evolution, like traditional methods, should be both flexible and conservative: they should preserve as much of the previous iteration as possible without leading to inconsistency. In addition, while many projects can aspire to *ontologies* that are useful for a wide array of other applications, dynamic ontologists can hope that their automated *methods* of ontology design will generalize as well.

Many of the automated metadata management tools available today operate primarily on term co-occurrence statistics. Term co-occurrence approaches attempt to recover semantic information about terms from the textual context in which they appear (whether it be sentence, paragraph, or entire document). As anyone who has used a search engine can attest, however, co-occurrence information alone is often not enough to intelligently infer semantic relevance. Even standard methods of augmenting co-occurrence methods—such as utilizing user searching and linking behavior, as in Google’s PageRank™ algorithm—do not reach the standards of reliability or transparency one desires in an academic reference work (Hinman 2005). The problem of automatically identifying semantic relevance is deep and abiding in computer science, and we do not expect a general solution which meets our reliability criteria anytime soon.

Rather than searching for a fully-automated solution to our metadata needs, we seek to utilize the SEP's most valuable resource—regular access to domain experts in philosophy. On our approach to dynamic ontology, we begin with a small amount of initial manual ontology construction. Once the initial structure is in place, a variety of automated methods are used to structure feedback solicitation forms and deploy that feedback in data validation, ontology population, and the semi-automatic extension of its taxonomic structure. Our hope is that by managing our access to domain experts as efficiently as possible, and by distributing feedback solicitation throughout the SEP's normal workflow, we can minimize or even eliminate the need for expensive double experts.

Once we have created an ontology for the discipline and populated it with individuals corresponding to SEP keywords (created a *knowledge base*), semantic relationships between terms can be read off of the ontology by humans or automatic agents through the taxonomic and non-taxonomic links it records between them, thus addressing the SEP's metadata needs.⁷

3. The Engineering Task

The engineering task facing our project is to efficiently and economically produce, populate, and maintain a viable dynamic ontology for the domain of philosophy. As such, we have created a process to semi-automatically generate a formal representation of the tools, products, attributes, and activities of the philosopher, with special emphasis on the category of philosophical ideas.⁸ The InPhO contains information about philosophical ideas and positions, an extensive array of biographical data about philosophers, citation information on the documents they read and produce, information about the organizations in which they participate, and much more.

[Figure 1 here]

Fig. 1: Protégé screen shot showing InPhO categories with sample instances.

3.1 Related Projects

With any technical project, it is worth reviewing other endeavors in the neighborhood. First, we note that we have been unable to locate any relevant ontologies of philosophical ideas in the standard ontology databases (Protégé databases, DAML, Ontolingua, Swoogle, for example); and, in general, there seems to be little work done on ontologies of ideas. Rather, most ontologies focus on more stable taxonomic structures (especially on types of physical objects or positions in social hierarchies), and few focus on the classification of abstract objects. This is to be expected, as the classificatory structure of abstract entities is much more likely to be unstable, vague, and controversial.

One relevant project is the Philosophy Family Tree⁹ maintained by Josh Dever. It uses genealogy software to record philosophers and their dissertation advisors—in some cases, all the way back to Leibniz. Dever has graciously allowed us to use his data to enhance our ontology, and we have incorporated its information into our current version.

The other project most directly relevant to our own is the PhiloSURFical project¹⁰ (Pasin & Motta, this issue). Pasin's research concerns the creation of an ontology which can be used to "describe philosophical resources, and allow an easy content-driven navigation of them." Pasin's approach is explicitly driven by pedagogical goals, especially the desire to formally annotate philosophical texts and take a learner on a "guided tour" of the text's logical structure. The current draft of Pasin's ontology, however, lacks the primary structure we require for our Information Retrieval (IR) needs: a decomposition of encyclopedia keywords along disciplinary lines, grouping together concepts and positions by mutual relevance. Our differing interests limit the degree of overlap while maximizing the possibility of future

productive collaborations. In particular, we are enticed by the idea of offering a learner a “guided tour” through the contents of the SEP as part of the SEP’s envisioned conceptual and semantic navigation interface.

The richest sources of information for our project have been two excellent annotated bibliographies maintained by SEP editors. David Chalmers has created a superb bibliographical taxonomy for the specialization of philosophy of mind.¹¹ In addition, there is an excellent philosophy of science bibliography maintained at the University of Pittsburgh (created by Rob Clifton, John Earman, and John Norton).¹² With the permission of the authors, both were incorporated with only small modifications into the ontology.

Finally, Wikipedia is increasingly adding more classificatory structure to its encyclopedia, and its various “table of contents” features contain extensive classification of philosophical ideas and positions. However, as is often observed, Wikipedia’s peer-supported nature—which brings with it the persistent possibility of inaccurate information, malicious manipulation of controversial entries, and vandalism—makes it unsuitable for academic referencing purposes.¹³ Our approach to Wikipedia’s metadata is to treat it as a valuable source of unverified input data, which is then passed to domain experts for validation before being incorporated into our representation. It is also worth noting that Larry Sanger, the estranged co-founder of Wikipedia, has announced the development of a new reference work called Citizendium,¹⁴ which adopts a much more expert-centered approach, and the new Scholarpedia¹⁵ is also trying to find ways of harnessing expert review to manage distributed authorship.

3.2 Guiding Principles

For those unfamiliar with the process of ontology design, it is worth noting three pieces of accepted wisdom that have been stated succinctly by Noy and McGinness (2000):

- 1) There is no one correct way to model a domain — there are always viable alternatives. The best solution almost always depends on the application that you have in mind and the extensions that you anticipate.
- 2) Ontology development is necessarily an iterative process.
- 3) Concepts in the ontology should be close to objects (physical or logical) and relationships in your domain of interest. These are most likely to be nouns (objects) or verbs (relationships) in sentences that describe your domain.

3.3 The Categories

Our ontology currently contains six basic categories: Thinker, Document, Organization, Nationality, Profession, and Idea. The “Thinker”, “Profession”, and “Nationality” sub-ontologies possess little classificatory structure and only contain as instances lists of philosophers, nationalities, and professions obtained by various sources (primarily by “slurping” data from Wikipedia and allowing authors/editors to manually insert missing individuals). In order to maximize data interoperability, top level concepts have been mapped into standard ontologies where possible. The Thinker category has been mapped into the W3C class FOAF::person (augmented with the additional biographical property slots).¹⁶ The hierarchical structure and properties of “Document” categories are adapted from the AKT Reference Ontology,¹⁷ and the “Organization” category was taken (with appropriate augmentation and pruning) from the Protégé ontology library¹⁸.

3.4 Dual Ontology for Ideas

Perhaps the most noteworthy aspect of our approach is its classification of ideas according to semantic inheritance relationships holding between the contents of ideas rather than more formal inheritance relationships observed in their types (e.g. social or structural roles). For example, the most straightforward ontological decomposition of the category *philosophical idea* might break it down into various social or structural roles—for example, *philosophical idea* → {*position*, *concept*, *problem*, *distinction*, *argument*, ...}, *problem* → {*dilemma*, *trilemma*,

paradox, ...}, *distinction* → {*bipartite distinction, tripartite distinction, continuum*}, and so on (see also Grenon & Smith, this issue).

However, this categorization—though useful for some purposes—would not answer to the metadata needs of the SEP. For one, this “social/structural” decomposition would likely be too shallow. We have several thousand concepts and need a rich classificatory structure which can separate them into meaningful clusters for the purposes of cross-referencing and semantic navigation. Moreover, we prefer a decomposition focusing on taxonomic structure that is as stable and non-controversial as possible, and thus would be familiar to most of the SEP’s authors and editors. No widely-accepted social/structural decomposition of philosophical ideas currently exists, so one would have to be engineered.

For these reasons, we chose to focus instead on a decomposition classifying ideas according to their locations in the semantic space of the discipline. Thus, our category *philosophical idea* breaks down into *idea about epistemology, idea about metaphysics, idea about ethics, idea about logic, idea about philosophy of mind*, and so on. (Note that hereafter, the “*idea_about...*” prefix is omitted from all category names, and should be implicitly assumed to avoid confusion.)

Information about the “idea type” is nonetheless useful for inferential purposes. For example, knowing whether an idea is of the type *position* or of the type *distinction* can constrain the types of relationships philosophers can have to it. The InPhO thus represents this information as non-taxonomic relations (e.g. ‘*is_idea_type(connectionism, position)*’).

3.4.1 Experts for finer structure

For the initial draft of our ontology, we directly solicited taxonomic schemes from editors of the SEP. Instead of imposing a single conceptual structure or set of design principles on all

subdisciplinary areas in the encyclopedia (for example, by trying to force general philosophical divisions like realism/anti-realism or subjectivism/objectivism onto political philosophy as well as metaphysics), the remainder of the idea category's initial build has been shaped by soliciting subdisciplinary taxonomies from area experts. By choosing not to normalize organizing principles across subdisciplines, we grant experts more freedom to provide their best conceptualization of the target domain. A benefit of this approach is that the information is current and likely to be highly semantically relevant, and most likely to lead to objective, comprehensive, and elegant representations of the subject matter (Sanger 2008). There are drawbacks, of course: individual experts are likely to be highly influenced by their own particular interests, and hence their taxonomic representation of ideas in a particular subject area may be disproportionately detailed in certain areas and overly sparse in others.

[Figure 2 here]

Fig. 2: View of InPhO interface showing subcategory structure.

Another important aspect of ontology design concerns the drawing of a line between categories and individuals. For instance, should *connectionism* be considered an instance of the category *philosophy of artificial intelligence*, or should it be treated as a subcategory with instances of its own? Such questions are always present in ontology design, but there can be no right answers without considering the intended applications. We have been guided by a rough rule of thumb that we should approximate a one-to-one correspondence between titles of SEP articles (unless they are articles about categories) and individuals in the ontology; thus, whether *connectionism* will be a category or individual will depend on the amount of treatment it receives in the current version of the SEP. It is also an important feature of our design that the ontology is revisable, and items treated as individuals in one iteration may be treated as categories in the

next. For example, though *connectionism* may initially be made an individual, as the encyclopedia adds more entries in cognitive science, it may become appropriate to treat it instead as a category with instances of its own such as *parallel processing*, *backpropagation*, *distributed representation*, and so on.

Another issue is that some concepts and individuals may not naturally belong in a single unique location in the ontology, being multiply classified by experts and/or automated methods. While most ontology languages permit multiple inheritance, rampant multiple classification erodes the classificatory utility of the ontology and increases the computational complexity of population (see next section).

We have attempted to finesse this issue by limiting concepts to a single appearance per subdiscipline while adding "semantic crosslinks" to the ontology. These crosslinks capture the relatedness of ideas deemed mutually relevant by feedback or automated methods yet which have been manually classified by experts in distant areas of the ontology. The inferred semantic information is important for editorial needs such as cross-referencing, as well as for readers who may wish to follow leads into related topics. Considering the reachability of concepts via semantic links also allows us to prune redundant arcs from the ontology and classify individuals in optimal locations during population (Niepert et al 2008). We believe that cross-links can be relatively sparse yet still provide these advantages, and additional computational complexity is minimized because our reasoning always begins with a significant amount of the ontological structure in place.

Such cross-links can compensate for the categorization of a concept by experts in a location which does not accurately reflect its importance in the SEP. For example, an idea which has been categorized by experts at a level of specificity which does not reflect its corresponding

term’s presence in the corpora will preserve its high degree of connectedness by having a large number of semantic crosslinks. For example, *philosophy of cognitive science* was placed under *philosophy of psychology* by a domain expert, but, given that the former phrase occurs in a wider range of SEP articles than the latter, our semi-automatic methods classified it at a higher level of generality and inferred many crosslinks to topics in other parts of the structure such as *theories of mental content*, *artificial intelligence*, *consciousness and intentionality*, and the *intentional stance*. Semantic crosslinks allow us to capture the relevance of *philosophy of cognitive science* to these otherwise distant nodes without directly contradicting the expert’s taxonomy.

[Figure 3]

Fig. 3: View of InPhO interface showing crosslinks (@) from philosophy of cognitive science to other topics. SEP thinker icons to the left of terms indicate links to Stanford Encyclopedia articles, while the ‘SEP’, ‘Noesis’, ‘Scholar’ links provide direct access to the search engines using a query that is composed of that term plus its superordinate category.

3.4.2 Properties (Non-Taxonomic Relations)

We sought to include all salient properties and relations in our ontology which could possibly be inferred by semi-automated means. In general, an ontology with more properties and relations is better than one with fewer, as it increases the scope of the knowledge base and the chances of bootstrapping via intra-ontology inference. Table 1 below contains a list of the initial property list for our ontology organized by domain and range.

Table 1: Initial Property List

<u>Thinker→Thinker Relations</u>	<u>Thinker Properties</u>
Teacher_of	Born_on
Influenced	Died_on
Criticized	Spoke_language

Defended
Dissertation_Advisor_of
Discoursed_with

Nationality (Thinker→Nationality)
Profession (Thinker→Profession)

Idea Properties

Idea_type (concept, position, etc.)

Thinker→Document Relations

Wrote
Edited

Thinker→Organization Relations

Member_of
Studied_at

Thinker→Idea Relations

Worked_on (problem)
Created_view
Attacked_view
Espoused_view
Aware_of

Document→Document Relations

Published_in (article → journal/book)

Idea→Idea Relations

Opposed_to
Commits_to (idea1 commits one to idea2)

Document→Idea Relations

Discusses

Ternary Relations

Disagreed_with(X, Y, Z): ThinkerX disagreed with ThinkerY on IdeaZ

3.5 Populating the InPhO Using Expert Feedback

Population (also called annotation) is the process by which a number of individuals are classified according to an ontology's taxonomic structure and values are supplied for those individuals' non-taxonomic properties ("slots"). Semi-automated population of taxonomic and non-taxonomic relations will take place through the solicitation of expert feedback as part of the SEP's document submission process. When SEP authors submit a document for publication, their feedback will be solicited in a three-stage process. First, statistical measures such as the widely used tf-idf measure (term frequency–inverse document frequency) and n -gram models (conditional probability within the corpus of a word given the previous n words) are run over their document to infer terms and names of possible significance which occur in the document; authors will be presented with a list of such terms and asked whether their document really discusses them or they occur in it only incidentally ("Is your article about *connectionism*?").

Second, the authors are asked to evaluate the relatedness (on a five-point scale from ‘totally unrelated’ to ‘highly related’) and generality (‘more general’, ‘less general’, ‘both’, and ‘incomparable’) of target terms they selected in the first stage to other terms highly ranked by our statistical methods as candidate hyponyms (of lesser generality) or hypernyms (of greater generality) to those target terms.

[See Figure 4 here]

Fig.4: The idea feedback interface.

Once we have obtained this feedback, the task is to use it to populate the ontology by classifying terms in its hierarchy which occur in the document but are not yet contained in the knowledge base. This is accomplished by representing the expert feedback as a series of facts in first-order predicate calculus¹⁹ and using non-monotonic inference techniques to infer the classifications induced by the feedback.²⁰ Currently, we use this method only to classify idea instances according to the existing ontology structure, but with appropriate feedback it could be extended to build (partially or completely) the idea ontology’s taxonomic structure itself. We intend to explore this extension, but we do not expect its results to be as reliable as the more modest method already implemented.²¹

Finally, experts are shown a third feedback page where they are asked to evaluate the non-taxonomic relations inferred by our automated methods (see Niepert et al. 2007). Statistical methods are used to infer relationships obtaining between terms occurring in the document. Author and editor feedback are used both to populate non-taxonomic relations in the knowledge base and to augment or construct training sets for statistical learning techniques such as Hidden Markov Models.

This three-step feedback-harnessing strategy demonstrates how we hope to achieve a higher degree of reliability than pure co-occurrence and wiki-style approaches. Our database will be separated into “clean” and “unclean” portions. The “unclean” portion consists of information inferred by taxonomic methods or gathered from other external sources (such as from Internet crawls and user search traces). The “clean” portion consists of information either gathered directly from experts (such as the taxonomic decompositions recovered from the annotated bibliographies) or approved by experts through the feedback-elicitation process. The threat of inaccuracy is addressed by assuring that only the clean parts of our database are used to infer the ontology and thus feature in cross-referencing, navigation, and other implementations viewable by the SEP’s users.

An advantage of our expert-supervised approach to ontology population is that it provides ample quality-control while still allowing for the possibility of novel discovery. Our co-occurrence statistics have turned up a number of connections that might easily have been overlooked by area experts. To take one example, the methods ranked *anaphor* as one of the highest hypernym candidates for *propositional attitudes*. Though two of the authors of this article are philosophers of mind, we initially thought this connection was due to error. A quick SEP search revealed, however, that the ranking could be explained by the fact that anaphoric sentences pose a challenge to the Fregean theory of propositional attitudes. This interesting connection would likely not be discovered by novices; and though experts might not think of the connection off the tops of their heads, they can easily uncover its validity with a cursory inspection of the relevant SEP articles.

A key challenge facing philosophers in the digital age is to discover how best to use computers to support our understanding of the discipline. As should be clear by now, we do not

expect our information retrieval tools to write a compelling philosophy paper any time soon.

Rather, we recognize what has long been true: that humans and computers work better together than either do in isolation. With careful collaboration, reliable representations of the discipline can be created facilitating a wide range of future tasks in digital philosophy (several discussed below). Keeping future tasks such as the design of visually-effective conceptual and thematic navigation tools in mind, we have tried not to “overfit” our methods and ontology to the needs of the editorial staff of the SEP.

3.6 Raw Materials—Exploitable Semantic Structure

In this section we describe six sources of exploitable semantic structure utilized in the course of our project. It may be read as a series of hints to any other projects seeking to design an ontology and automated IR tools for their reference works.

1. The SEP’s Editorial Structure: The richest source of large-scale classificatory structure available to us was the editorial structure of the SEP. Editorial oversight of the encyclopedia is divided into twenty subject areas corresponding to widely-recognized academic subdisciplines of philosophy (*logic, ethics, metaphysics, philosophy of action*, and so on). Within most of these subject areas, each article is under the jurisdiction of one subject editor, with the articles in a given subject area subdivided between several subject editors. We could infer a considerable portion of the large-scale classification of ideas by noting the editorial jurisdiction under which the articles about those concepts fell. (It is important to note, however, that the editorial structure of the SEP is partly a matter of administrative convenience and contingency.)
2. Current Cross-References: Another rich source of information on semantic relevance is provided by the current cross-references of articles, which have all been hand-coded by

editors and contributors. Such information can be taken as objective feedback and incorporated into future, semi-automated cross-references and used to instruct future sets of cross-references as they are dynamically updated.

3. Article Citations: The citations of SEP articles can also be used to infer relevance of articles (and thereby derivatively to the concepts the articles are about) under the assumption that semantically-related articles will tend to cite the same influential papers.
4. Article Titles: The article titles were the biggest gift provided by the structure of the SEP. Article titles, by unspoken convention, have largely been kept spartan and neatly-correlated with concepts or philosophers that the article is about. This makes it easy to build a one-to-one correspondence between individuals in the ontology and articles in the SEP.

Though our current draft of the InPhO only operates at a level of specificity corresponding to that of subjects of whole articles, we eventually hope to incorporate more specific ideas and terms by moving to the level of article sub-sections or even paragraphs. If we do so, section sub-headings will also be very useful in recovering more specific information about the themes of passages.

5. User Search Traces: The final useful source of information to be explored is user search traces. This information becomes significant on the assumption that users tend to search for semantically-relevant topics in the same session, although there are reasons for skepticism about this assumption. Nevertheless, we believe that looking for convergence of searching behavior among different users can also be used to improve our cross-referencing system.

4. Challenges

The following section presents a summary of challenges we encountered in the course of designing and implementing our ontology, which may also confront similar projects. Each section contains a summary of the difficulty, presents our current solution, and briefly theorizes about the origin and scope of the difficulty.

4.1 Problematic Category Names

In constructing the ontology, we found that many seemingly straightforward category names provided by experts were difficult to adapt to the constraints of our statistical methods (which operate using semantically significant keywords and their co-occurrence statistics).

Below is a list of types of category names which proved troublesome.

4.1.1 “Grab Bag” Category Names

Examples: “Other Psychophysical Theories,” “General Issues in Philosophy of Science,” “Topics in Feminism”

The most common problematic category names were those which seemed to be “grab bag” categories. A “grab bag” category here is defined as one which groups together possibly heterogeneous elements which are “leftovers” from more significant sibling decompositions; in other words, the category name itself occurs rarely (if at all) in the SEP, and its significance derives from a mutual exclusion relation to its more significant siblings. This sort of problem arose throughout the ontology, but we may focus on a useful case study which emerged in the *metaphysics of mind*.

In one expert-solicited taxonomy for the philosophy of mind, for example, *metaphysics of mind* decomposes into the subcategories {*materialism and dualism*, *functionalism*, and *other psychophysical theories*}. This decomposition makes intuitive sense; views on materialism and dualism form one closely-related way to conceive of issues in the metaphysics of mind, various ideas related to functionalism comprise another broad way to tackle the subject, and finally there

are other less popular but still significant psychophysical views which do not clearly fall into the other two categories. The name “other psychophysical theories,” however, is somewhat useless to methods which rely on co-occurrence statistics.

The frequency of such categories in the taxonomies supplied by experts suggests to us that, rather than simply demonstrating a lack of a label due to ‘laziness’, they seem to be doing some real classificatory work. Unfortunately, it isn’t clear how to replace these categories with something more informative and statistically tractable, even when they are considered on a case-by-case basis. The problem is not crippling for our methods where the grab bag category decomposes into more meaningful subcategories; where they do not, however, we have excluded them from the current build of the ontology.²²

4.1.2 “Philosophy of X” vs. “X”

Examples: “Philosophy of Connectionism” vs. “Connectionism”, “Philosophy of Cognitive Science” vs. “Cognitive Science,” etc.

In numerous places in the expert-supplied taxonomies, we also noticed that category names sometimes included a “philosophy of” prefix and sometimes did not; some examples include “artificial intelligence” vs. “philosophy of artificial intelligence,” “connectionism” vs. “philosophy of connectionism,” and so on. As knowledge modelers, we then must ask ourselves: does each member of the pair of labels denote the same idea? In the context of philosophy and as far as the co-occurrence methods are concerned, the answer frequently seems to be “yes.” When we determined this to be the case, the “philosophy of” prefix was discarded, but not when it was obvious that the two denoted robustly different entities (e.g., “philosophy of logic” vs. “logic”).

4.1.3 Category Names Containing Prepositions: “and”, “in”, “of”

Examples: “Language and Society,” “Reference and Denotation,” “Materialism and Dualism,” “Consciousness and Qualia,” “Consciousness in Physics,” “Cosmology of Physics”

Again, our methods must interpret the contents of the ontology in terms of keywords and co-occurrence statistics. Where possible, therefore, it is best to translate category names involving keywords joined by prepositions into keywords joined by union or intersection; in other words, for each category name of the form “Keyword1 <preposition> Keyword2,” we should decide whether both Keyword1 AND Keyword2 should be relevant to the contents of the category, or whether either Keyword1 OR Keyword2 may be relevant with or without the other term.

In general, “in” and “of” signify contexts in which both terms should be relevant, and this is represented by the intersection of the categories of ideas about each keyword; we use automatic scripts to replace these with the intersection symbol. “And,” however, was ambiguous throughout the ontology, and no principled translation rules could be developed without surveying the contents placed in that category by the experts. “(Ideas about) Materialism and Dualism” for instance, expresses the union of the category of ideas about materialism with the category of ideas about dualism, whereas “(Ideas about) Consciousness and Physics” expresses the category of ideas in the intersection of the categories of ideas about each keyword. There seems to be no way to mitigate this ambiguity at this stage of our project without having a human operator manually disambiguate after surveying the contents of the category.

4.1.4 Subconcepts involving anaphor or omitted adjectives

Example: “Mental content → theories of content”

Such decompositions are easily intelligible to humans, but it is difficult to automatically resolve (or even notice) the anaphor in the above example. Without resolving this problem, the automated methods may conclude that the children of “theories of content” are relevant not just to “mental content”, but also to linguistic content, intentional content, teleological content, and

so on. The only solution to this problem seems to be to prohibit such anaphor in the ontology and carefully ferret out any possible anaphoric or ambiguous category names.

4.2 Change and Evolution of the Ontology

As our ontology is populated over time, we may predict that the manually-coded classificatory structure will become increasingly obsolete. In particular, as some areas of philosophy become increasingly neglected and others become more fertile, some categories may shrink (as, for example, articles are deleted or radically edited) while growth in others (as a result of new articles being commissioned and added) may result in clusters which are larger than optimal for our metadata needs. Even more drastic changes to the idea category may be required by intellectual sea changes, such as those caused by the advent of new technologies (e.g. computers), the discovery of new scientific theories (e.g. quantum mechanics), or the growth or decline of broad philosophical movements (e.g. logical positivism).

There seems to be no way around the conclusion that large-scale paradigm shifts in philosophy will require fresh re-conceptualizations of the field from the top down. Rather than being fatalistic, however, we can note the advantages dynamic ontology has over other methods. First, re-conceptualizations which result from large-scale paradigm shifts may not have to be coded entirely from scratch; much of the old ontology may be reincorporated or used as inspiration where relevant, and metrics about the unsuitability of the obsolete ontology may guide the creation of the new one. Second, smaller changes to the ontology are tractable by modest semi-automatic means. For example, we will be able to automatically detect when some categories become too large, and then use automatic methods to divide the contents of the encyclopedia into two or three more closely-related clusters; human users may then be prompted to provide category names for these newly-created clusters by surveying their automatically-

populated contents. Similarly, categories which become defunct as a result of article deletion or radical editing of articles may be detected and pruned from the ontology.

A further advantage of managing change in any formal setting is the possibility of comprehensively tracking changes within the discipline over time. Versions of the ontology can be saved and archived every time a change is made to the ontology, and intellectual trends could be studied with precise metrics. This should be an enticing prospect for academics in a number of different areas interested in the change of ideas and socially-accepted conceptualizations, from metaphilosophy to history to social network theory.

4.3 Author/Editor Compliance

Another worry centers on the issue of author/editor compliance: why should volunteer authors and editors be willing to provide all the necessary feedback? Though we admit that this is a challenge, we feel that numerous incentives encourage SEP authors to supply feedback. First, we can note that the metadata production is currently completed manually in a time-consuming manner—authors or editors provide cross-references using the labor-intensive and imperfect method of searching through the SEP. By contrast, the feedback form conveniently presents authors with a thorough list of candidate cross-references ranked by relevance.²³ Second, the authors are submitting their articles as a matter of professional service and to increase their profile as an expert on a given issue; the augmentation of their article's metadata (and of the success of the SEP in general) increases the accessibility and visibility of their articles. Finally, we have designed the feedback forms to be as ergonomic as possible, and the amount of feedback we need from any one author is minimal—thus the process should not place undue demands on any one author's time and attention.

5. The Pies in the Sky

5.1 Semantic and Thematic Navigation

Perhaps the most important way the structure of the InPhO might be used to make the SEP's content more accessible involves the creation of a system for semantic and conceptual navigation. In such a system, an information visualization interface would allow browsers of the SEP to navigate by visually following semantic relationships between the SEP's ideas, thinkers, documents, and organizations. An envisioned user interface with the system is described as follows: A user will search for some specific keyword; a graph will then be displayed with the closest keyword contained by the ontology displayed in the center; the graph will be populated with a number of other related keywords linked to the search keyword, where the links are labeled with the relationship recorded in the ontology; the user can browse this structure, re-center it, and repopulate it; and each time the user clicks on one of the concept bubbles, a small window will be displayed showing the first few lines of the SEP article most correlated with that keyword. Moreover, the user can control the type of links displayed by checking various options, such as focusing on philosopher→philosopher relationships, idea→idea relationships, etc. (a mock-up of this system will make these ideas clearer).

[Figure 5 here]

Fig. 5: Mock up of navigation interface using ontology structure and contents.

Such a system would be an excellent tool for a reference work; it would facilitate interaction with the encyclopedia for expert and novice users alike. Experts could see the “big picture” of an idea or philosopher in a glance, capturing a large amount of information in an efficient manner, and novices could surf through the encyclopedia's content by following paths hewn by an expert-level understanding of the domain. Students engaged in research could instantly see which other topics are importantly related to their current search item, and the

semantic labels of those links would tell them not only *which* other items are related to their current search but also *how*.

We have made some exploratory progress in this direction using the java-based visualization package JUNG.²⁴ While this package comes with an impressive set of customization features, we would prefer to eventually develop our own visualization package designed especially for ontology-guided navigation of dynamic reference works (Crampes and Ranwez, 2000). Plans for such a system have been developed, and we are pursuing further funding opportunities in order to implement them. We believe that just as scientific visualization has become an important area of research for the sciences, philosophy visualization should be an important topic for philosophers. What can be visualized? And what visual metaphors guide our comprehension of philosophy? These are questions whose answers affect the development of philosophy as a discipline in the digital age.

5.2 Exploration of the Amount of Semantic Structure Retrievable from Co-occurrence Data

How much semantic information about terms is recoverable from co-occurrence information alone? How much, instead, can only be recovered by agents possessing human perceptual apparatus, physical or causal interaction with the subject domain, memory, emotions, human cognitive dispositions, neural architecture, and so on? Despite great advances in the last few years on both co-occurrence models and other models of semantic structure, and on language learning and semantic memory, no general answer to this question is currently known.

One advantage of our project is that it would provide a means by which this question could be explored in the domain of philosophy. As our IR methods are implemented and tweaked over time, we may be able draw generalizations as to which items or areas which they handle comparatively well and which comparatively poorly. Moreover, the ontology and the

SEP's corpus would provide a playground for new IR methods and cognitive models to explore these questions in a live domain.

5.3 Inferring Large Sections of the InPhO Dynamically

Thus far, we have discussed two ways in which the InPhO is a dynamic ontology: the dynamic population of instances and the dynamic extension of the ontology's category structure when sections become too large. An alluring future direction which could make the InPhO even more dynamic involves the use of more ambitious automated reasoning techniques to infer large sections of the ontology using automated means.

There are two distinct but related approaches here we would like to pursue. The first involves "relaxing" some of the rules in our answer set programs such that instead of producing a single ontology, they produce a large space of possible ontologies which are consistent with a general theory of ontologies (satisfying acyclicity and other constraints) and current feedback. A metric would then be applied to rank these possible ontologies according to their suitability, likely using a mathematical measure of the degree to which each ontology fits the predictions of relatedness and generality derived from the co-occurrence statistics. Another approach uses Markov Logic Networks (Richardson and Domingos, 2006), which combine first-order knowledge bases with probabilistic (Bayesian) networks. This method would provide a probabilistic ranking of the populated ontologies through the assignment of weights obtained directly from the statistical generality and relevance scores. A significant difference between these two approaches is that in the latter the populated ontologies are ranked directly as part of its inductive reasoning method, while in the former the rankings are computed as an independent, post-processing step. We hope to compare the two methods and determine which best serves the needs of dynamic reference works.

6. Philosophical Issues and Doubts

Before concluding, we move to address five worries and objections we have encountered to our approach.

6.1 What is the relationship between the InPhO and Metaphilosophy?

Some readers may be skeptical of our statistical approach, supposing that we are engaged in a very dubious metaphilosophical methodology. We want to make clear that we do not take the InPhO to show how philosophy *ought* to be organized. Such an ideal conceptualization is neither recoverable by currently-available automated means nor required for our purposes. Even human experts in metaphilosophy would be hard-pressed to produce a comprehensive decomposition of the conceptual space of philosophical ideas, and no doubt any such scheme would be the subject of controversy. However, we have real and pressing information management needs, and, for the purpose of meeting those needs, having an imperfect formal representation of the structure of the field is better than having none at all.

Despite this caveat, we argue that the InPhO can play a non-trivial role in metaphilosophy. A common sort of distinction may be useful here in explaining this role. Let us call a metaphilosophical position *descriptively adequate* if it portrays the discipline as it is currently practiced in such a way as to support practical applications. Let us call a metaphilosophical position *normatively adequate* if it outlines the best way to arrive at philosophical truth, provides philosophy with a firm epistemological grounding, or otherwise describes how philosophy *ought* to be organized to maximize our chances of progress. While we do not defend the InPhO's normative adequacy, we believe that the collaborative, distributed, and empirical approach used to construct it makes it more likely than alternatives to produce a representation of the discipline which is descriptively adequate.

Traditional methods for obtaining a description of the intellectual space of the discipline take the form of a less than a handful of philosophers testing each other's intuitions about ways to organize the content of the discipline. There are many reasons to be skeptical about this approach. First, there is the matter of individual bias; philosophers are likely to feel that the particular issues they work on are the most important. Second, the learning histories and intellectual trajectories which shape one's metaphilosophical standpoint are idiosyncratic and politically charged. Third, the drive for simplicity and elegance is a double-edged sword; it may tempt philosophers to artificially impose a normalized structure on distant areas of philosophy. Fourth, once settled on an approach or a set of organizing principles, overconfidence and confirmation biases may set in, leading the philosopher to feel the chosen approach is appropriate in diverse settings. And finally, given the explosion in the number of significant philosophical publications mentioned in the introduction, it is unlikely that even a reasonably-sized subset of extremely well-read philosophers would be sufficiently qualified in all areas of philosophy to produce a thorough conceptualization.²⁵

Viewed against these challenges, our approach may be thought to possess significant empirical advantages over this traditional method. Granted, our reliance on expert conceptualizations for the InPhO's basic framework renders us vulnerable to some of these worries—but we do not ask experts to represent anything outside of their own area of specialization. Furthermore, our feedback system is deployed with redundancy in mind. We propose the same generality and relevance hypotheses to multiple experts in the feedback process, and the nonmonotonicity of our logic programming methods allows us to flexibly respond to expert disagreement (Niepert et. al 2008). Our hope is that this will afford the InPhO a degree of intersubjectivity. We also note that basing our relevance and generality estimates on

the actual text of the SEP provides an independent and perhaps more objective window on the structure of the discipline than does the impressions of the authors and editors of those articles. We record not just what philosophers *think* they are writing about, but also study statistical properties of their *actual output*. For example, an author may think the central ideas of their article are X and Y, but if they spend five pages dealing with objections pertaining to Z (that the author would perhaps prefer to ignore), this is significant information that should be taken into account when deciding what the text is actually about. Thus while it remains to be proven that our method produces a more accurate descriptive metaphilosophical framework, there are a number of good reasons to suppose it will.

Furthermore, while we agree that there is little reason to suppose that the InPhO presents a normatively adequate metaphilosophy, the InPhO may be useful in the search for one. While foundationalist approaches to metaphilosophy may recommend sweeping away the debris and starting again from first principles, iterative approaches require a descriptively adequate picture of the current state of affairs. And if our goal in metaphilosophy is to assess the suitability of the current way of conceptualizing the intellectual space of the discipline (which includes assigning relative importance to various ideas, positions, problems, and areas of philosophy), a formal representation may prove more amenable to analysis than an informal one.

A last word on this issue: while we do not defend the InPhO's normative adequacy, we do retain a keen interest in the suitability of our ontology's specific structure and conceptualization scheme, especially as it is evaluated by its users (novice and expert alike). We welcome further additions to the metacontent structure we have inferred, including extensions of the existing ontology as well as alternative, mutually inconsistent classification schemes. Our metadata management and visualization methods are likely to be more effective with the help of

multiple ontologies, and so we welcome suggestions—and even contributions of alternative ontologies from readers who find themselves dissatisfied with ours.

6.2 Is the InPhO properly called an ‘ontology’?

In presenting this material, several critics have lamented our terminology. This criticism has come from at least two different directions. The first line of criticism comes from those who believe the word ‘ontology’ has been stretched too far from its original metaphysical roots, feeling that it is more misleading than useful when referring to a formal data representation rather than an actual hierarchy of being. This group of critics—which has included at least one lexically-conscious information scientist, as well as philosophers—tends to object to the use the term has acquired in the computer and information sciences in general.

Another, more focused line of objection has come from those who prefer to reserve the use of the term for a particular kind of a formal representation used in the overarching project of the semantic web. These critics tend to focus on our idea category, indicating that the kind of semantic taxonomy expressed there is somehow not properly described as ontological. There are a number of things one might note here: it isn’t clear how the semantic divisions we focus on will enhance automated reasoning, or how the category of philosophical ideas will plug into higher-level ontologies to ensure interoperability (e.g., SUMO, Pease et. al 2002). Or, relatedly, one might worry about the social and dynamic nature of the conceptualizations we elicit through the feedback process, as domain ontologies in the semantic web are supposed to be authoritative descriptions of the types of entities in a domain suitable for a wide array of purposes (Smith 2003). The most serious form of this last worry centers on the semantics of the *isa* relation. The *isa* class subsumption relation is often presumed to hold in a context-general way and capture “purely ontological” relations; if it is true that “A *isa* B”, then any A in any context should by its

nature be a B. Because our idea category is organized primarily by contingent semantic relationships holding only in the context of certain philosophical discussions, its *isa* links may fail to hold domain-generally and may conflate taxonomic and non-taxonomic relations.²⁶ For example, the set of “ideas about connectionism” may be a subset of “ideas about philosophy of mind” in the context of philosophy, but not in the context of computer science or network theory.

We concede this last point; the classifications under the idea category will likely not hold domain-generally. For these divisions, we propose that the *isa* relation be understood to express *conditional* subsumption: A *isa* B, but only in the context of C, ..., Z, B’s ancestors. Let us call this conditional subsumption relationship ‘*isa**’.²⁷ Rather than supposing that this prevents the InPhO from being a proper ontology, we think the move from *isa* to *isa** helps capture precisely the sort of information required for our metadata needs. We believe that this sort of conditional, hierarchically-structured knowledge is important in modeling the ability of philosophical experts to say which topics are most relevant to the examination of particular ideas in particular scholarly contexts.

One might worry that the *isa** relation will have a mysterious semantics or that it might harm data interoperability. While dealing with *isa** relations will be more involved than dealing with good old *isa*, the previous paragraph gives the relation a straightforward technical semantics. As for determining data interoperability, the simplest schemes rely on matching types to one another across diverse representations in a one-to-one correspondence, often with the help of upper-level ontologies; let us call this “context-independent interoperability”. More sophisticated methods for determining interoperability are possible, however, which might for example compare contextual features of the two data representations (in this case, the two candidate types’ ancestors) to determine the their compatibility. For example, we may

determine that data are not interoperable between two *isa** ontologies if “philosophy of cognitive science” is a child of “philosophy of mind” in one ontology and “philosophy of science” in another. Finally, we note that only the idea category of the InPhO is structured around *isa**. The rest of the data in our knowledge base should have context-independent interoperability, and all of the subcategories of the idea category can be discarded in the process of translation if needed. In other words, the thousands of idea keywords can simply be treated as instances of the class *philosophical idea*, and their properties and relationships to instances of other categories should be as close to the ideal of context-independence as is realized in other standardized ontologies.

While these observations clarify the issue, it may remain a point of dispute whether a representation based on the *isa** relation is properly called an ontology. Our general response to the terminological worries of this section is to acknowledge the need to regiment language for purposes of clarity and precision, but note that all we mean by “ontology” and “dynamic ontology” is precisely what we say in Section 2 above. To the second line of response, we can also reply that the type of authoritative, domain-general, “final answer” description often sought by semantic web researchers would not prove responsive to the SEP’s full range of metadata needs (for the reasons why, the reader may glance back at the discussion of the “social/structural” option in Section 3.4). Though we would be reluctant to give up the word “ontology” entirely, readers sufficiently bothered by this issue may call the idea section of the InPhO what they like—we place no special importance on its terminological status.

6.3 Stagnation, self-confirmation, and bias

One of the general doubts we have faced when explaining our project goes something like this (to preempt critics from seizing upon a tempting pun): “You aren’t engaged in a *taxonomy* of ideas, but rather a *taxidermy* of ideas. Philosophy is by its nature a fluid, ongoing

conversation with rules always in flux, and to attempt to tie it to some conceptualization scheme or another is to destroy the essence of dialectic inquiry, reducing it to something sterile and irrelevant.” In response to this concern, we can agree that philosophy is in flux but deny that attempts to represent the current state of this flux are necessarily inadequate or stultifying to the discipline. Encyclopedias have not had this effect, because philosophers and their students have always regarded what they contain as subject to challenge and revision. There is no reason to expect that other attempts to represent philosophy will be treated any differently. Furthermore, the dynamic methods used to generate, populate, and evolve the InPhO over time should render it reasonably responsive to the ongoing dialectic of the discipline.

Secondly, one may argue that there is a bias in either the content of the SEP, the design principles of the ontology, or both. For example, one might accuse the SEP or InPhO of a pro-scientific bias, an anti-Continental bias, a pro-Western bias, a sexist bias, and so on. Furthermore, one might think that not only is the SEP and/or the metaphilosophical picture expressed in the InPhO biased or inaccurate, but giving it wide exposure is also likely to only exacerbate the problem. For instance, once the ontology has been deployed and comes to influence the way that cross-references, tables of contents, and navigational tools are designed and used in a widely-viewed reference work, it makes it more likely that “metacontent inertia” will set in, strengthening the current biases of the SEP in future versions. Moreover, it seems that most of the methods proposed to change and evaluate the ontology are not independent of interaction with that ontology; experts may be influenced by their constant interactions with it, and users who learn from interaction with the visualization tools are likely to accept its conclusions and confirm them in future sessions (as expressed in search traces).

Our central response to these worries is that we are only aiming to model the way that the discipline (as practiced by certain experts in the roughly Western, analytic tradition) sees itself. If there is a bias in this viewpoint or in the SEP itself, we think that formally modeling it may actually shine a light on it and spur reform. Furthermore, the act of constructing an ontology using data drawn from the SEP can help reveal imbalances in its coverage, or areas that are not sufficiently sub-divided into specialized topics. Analysis of the statistical properties of the entries and the generality of the terms they contain can help ensure uniformity of treatment across the discipline.

After all, it isn't like philosophers to allow standards or accepted truths to stifle such progress; in fact, the history of philosophy shows that stating something precisely and making it widely available is instead an invitation to more intelligent critique. Given the contentious nature of philosophy and its practitioners, therefore, it seems more likely to these authors that codifying philosophy's self-conceptualizations will *enhance* metaphilosophical critique rather than nullify it.

7. Conclusion

In brief, just as human understanding of everything, from geometry to the human genome, has been enhanced by developing new ways of representing it, so the discipline of philosophy also stands to gain much in being represented in explicit, formal, and computationally-tractable ways. We have argued here for the advantages of a dynamic approach to representing the discipline of philosophy. The near future of representing the discipline proves to be exciting, as many other approaches are possible, and it is largely an empirical question which will turn out to have the best combination of advantages and disadvantages for particular applications.

¹ For a discussion of the impact of this explosion on the discipline of philosophy in particular, see Rescher (2007).

Rescher believes that one effect the explosion has had on the discipline of philosophy has been to reduce the relevance of the “Great Man” approach to its history—to hold that the agenda of philosophy is determined by a “dominant elite” and that if one follows the work of the several greats of the time, one captures everything of import. Rather, Rescher thinks we should instead recognize that “philosophical innovation today is generally not the response to the preponderant effort of pace-setting individuals but a genuinely collective effort that is best characterized in statistical terms.”

² For an overview of the ineffectiveness of search, see for example Holscher and Strube (2000). Their unsurprising finding is that the initial searches of novices and subsequent corrections of their search string are largely ineffective. More interestingly, their research also suggests that neither expertise in the content domain nor expertise in web search are alone sufficient to produce a high success rate on common searching tasks; rather, a high success rate is only achieved by users with both high web experience and high domain knowledge. Thus the paradox of accessibility for digital reference works: those users most in need of the information that encyclopedias offer are the least able to find it.

³ <http://inpho.cogs.indiana.edu/>

⁴ <http://plato.stanford.edu/>

⁵ For example, the InPhO is to be integrated with the Noesis philosophy search engine (<http://noesis.evansville.edu/>) and the InPhO interface itself already provides for ontology-guided search in the SEP, Noesis, and Google Scholar (see Figure 4).

⁶ Put more precisely, we take a computational ontology to be a directed acyclic graph where nodes represent concepts and the links between concepts represent the taxonomic “is_a” relation (e.g. in the graph where ‘*Wine* → *Red Wine* → *Tawny Port*’, everything that “is a” instance of *Tawny Port* “is a” instance of *Red Wine*, and everything that “is a” instance of *Red Wine* “is a” instance of *Wine* (we are less strict as to whether this subsumption relation must hold in all contexts of discourse—see section 6.2). A knowledge base is an ontology that has been populated with individuals; mathematically, knowledge bases contain another kind of link denoting the “instance of” taxonomic relation, and a new kind of node denoting individuals, (with the restriction that no individual can have any children). Thus we may populate our toy ontology with an individual, ‘Fladgate Tawny Port’, which is an instance of the concept *Tawny Port*. A computational ontology will also contain declarations for a number of non-

taxonomic relations, which can either hold between individuals and constants (properties) or between two (or more) individuals (relations). Instances of these non-taxonomic relations may also be encoded in the process of ontology population. For example, a knowledge base might record the property ‘alcohol_content(Fladgate Tawny Port, 20%)’ and the relation ‘pairs_with(Fladgate Tawny Port, brie)’. Finally, a computational ontology may also contain a number of inference rules and axioms that can be used to reason about objects in the domain (e.g. infer the presence of certain relations on the basis of others, enforce consistency and default properties, and so forth).

⁷ This approach has a further semantic advantage over mere co-occurrence approaches in that the semantic relationships come marked with taxonomic and non-taxonomic types. The semantic types here will include taxonomic relations (e.g. “more general”, “less general”) and non-taxonomic relations (e.g. “teacher of”, “wrote”, “defended”, “had_profession”). There are several additional reasons why one might desire an ontological approach to the problem of similarity noted in Noy and McGuinness (2000). For one, it allows domain assumptions to be made explicit and rendered in a form more conducive to analysis and critique. This may be especially useful for a large encyclopedias like the SEP—for if the ontology is generated through a combination of co-occurrence methods and author feedback, then any bias observed explicitly in the ontology is likely present implicitly in the SEP itself. Making such biases explicit (especially overemphasis or underemphasis) provides SEP editors with useful information when commissioning future entries. Secondly, coding the information in a formal language like OWL (Web Ontology Language) allows it to be reused and exported to other projects and purposes. Finally, having a comprehensive, machine-readable, open-access knowledge base about philosophers and philosophical ideas is of interest in its own right.

⁸ The scope of “philosophers” here ranges over professional academic philosophers in the Western tradition.

⁹ <https://webspace.utexas.edu/deverj/personal/philtree/philtree.html>

¹⁰ <http://philosurfical.open.ac.uk/index.html>, and see also Pasin et al (2007) and his essay in this issue of *Synthese*.

¹¹ <http://consc.net/biblio.html>. David Chalmers and David Bourget are currently developing an extension of this taxonomy to cover all areas of philosophy, which will be called Philpapers (see <http://consc.net/taxonomy.html>).

¹² <http://philsci-archive.pitt.edu/view/subjects/>

¹³ See, for example, Jaschik (2007). But see also Giles (2005), who argued that Wikipedia, while error-prone, is not much worse off than *Encyclopedia Britannica*. In this debate, it is important to note that errors in Wikipedia are often self-promotion or intentional vandalism and are thus often more fantastic and malicious than those of

traditional encyclopedias like *Britannica*. Finally, Wikipedia's lack of archiving features for citation purposes is also problematic. Wikipedia should perhaps not be faulted for these shortcomings; while it is a goal of the SEP to be academically-citeable, it is not clear that it is ever a good idea to cite a domain-general encyclopedia (like Britannica) for anything but a survey of common opinion.

¹⁴ See <http://www.citizendium.org/>. See also Sanger (2008).

¹⁵ <http://www.scholarpedia.org/>

¹⁶ http://xmlns.com/foaf/spec/#term_Person

¹⁷ <http://www.aktors.org/publications/ontology/>

¹⁸ http://protegewiki.stanford.edu/index.php/Protege_Ontology_Library

¹⁹ For example, if the author indicates that *connectionism* is highly related and more specific than *philosophy of mind*, we represent these facts as “highly_related(*connectionism*, *philosophy of mind*)” and “more_specific(*connectionism*, *philosophy of mind*)”.

²⁰ This task is described and advantages of our approach discussed in full detail in Niepert et al (2008).

²¹ See Niepert et al (2008) for ideas about how this could be accomplished.

²² We are interested in exploring means to help the automated methods manipulate the category name in the same methods that humans do, perhaps by a kind of exclusion rule with its sibling categories. In this case, then, we might make “Other Psychophysical Theories” significant to the IR methods by translating it to a category comprised of the keywords “Psychophysical,” “Theory”, and “Metaphysics of Mind” and exclude the keywords “Materialism”, “dualism”, and “functionalism.”

²³ This provides an advantage over the approach of Kim et al. (2007) which provides technological facilitation for human construction of a philosophy ontology, but does not automate the processes for discovering the semantically significant relationships among philosophical objects.

²⁴ <http://jung.sourceforge.net/>. For an analysis of JUNG's capabilities, see O'Madadhain et al (2005).

²⁵ We might look, for example, at the top-level classification scheme of the Routledge Encyclopedia of Philosophy <http://www.rep.routledge.com/signpost-articles>. One may object to philosophy of psychology falling under philosophy of mind rather than (also) philosophy of science; one may object to the article on African Religion coming under the heading of African Philosophy rather than World Religions (showing a Eurasian bias). Our purpose here is not to point any fingers at the REP, for similar concerns could be raised about classifications in the

InPhO; we would, however, like to raise a question about how such disputes are best settled. We believe that they are better settled by looking at statistical properties of philosophical texts and soliciting feedback from multiple experts rather than by asking a one or two experts their opinion of the best classification.

²⁶ For related problems faced by projects attempting to convert the Wordnet's hypernym/hyponym taxonomy into an ontology, see Trautwein & Grenon (2004) and Oltramari et al (2002).

²⁷ For a related approach, see Stuckenschmidt (2006).

Acknowledgments

This research has been funded by Indiana University under the grant “New Frontiers in the Arts and Humanities” and by Digital Humanities Start-up Grant HD-50203-07 from the National Endowment for the Humanities. We also wish to thank the members of the Experimental Epistemology Laboratory for helpful comments on earlier versions of this document.

References

- Ceusters, W. & Smith, B. (2006b). A realism-based approach to the evolution of biomedical ontologies. *Proceedings of the Annual AMIA Symposium*, 121–125
- Crampes, M. & Ranwez, S. (2000). Ontology-supported and ontology-driven conceptual navigation on the world wide web. *Proceedings of the 11th ACM conference on Hypertext and Hypermedia*, 191-199
- Flouris, G., Plexousakis, D., & Antoniou, G. (2006). Evolving ontology evolution. *Proceedings of the 32nd International Conference on Current Trends in Theory and Practice of Computer Science*, 14-29
- Giles, J. (2005). Internet encyclopedias go head to head. *Nature*, 438, 900-901
- Gruber, T. (1993). A Translation Approach to Portable Ontology Specifications. *Knowledge Acquisition*, 5(2), 199-220

Gruber, T. (2008). Ontology. (In *Encyclopedia of Database Systems*, Ling Liu and M. Tamer Özsu (Eds.) *Springer-Verlag*.) (in press)

Guarino, N. (1995). Formal ontology, conceptual analysis, and knowledge representation. *International Journal of Human-Computer Studies*, 43, 625-640

Hinman L. (2005). Esse est indicato in Google: Ethical and political issues in search engines. *International Review of Information Ethics*, 3(6), 19-25

Holscher C. & Strube, G. (2000). Web search behavior of internet experts and newbies. *Computer Networks*, 33, 337-346

Jaschik, S. (2007). A stand against Wikipedia. *Inside Higher Ed*. Retrieved July 7, 2008 from <http://www.insidehighered.com/news/2007/01/26/wiki>

Kim, J., Choi, B., Shin, H., & Kim, H. (2007). A methodology for constructing of philosophy ontology based on philosophical texts. *Computer Standards & Interfaces*, 29, 302–315

Niepert, M., Buckner, C., & Allen, C. (2007). A dynamic ontology for a dynamic reference work. *Proceedings of JCDL 2007*, 288–297

Niepert, M., Buckner, C., & Allen, C. (2008). Answer set programming on expert feedback to populate and extend dynamic ontologies. *Proceedings of 21st FLAIRS*

Noy, N. & McGuinness, D. (2001). Ontology Development 101: A Guide to Creating your First Ontology. *Stanford Medical Informatics Technical Report 880*. Retrieved January 8, 2009 from <http://smi.stanford.edu/smi-web/reports/SMI-2001-0880.pdf>

Oltramari, A., Gangemi, A., Guarino, N., & Masolo, C. (2002). "Restructuring Wordnet's top-level: The Ontoclean approach." Paper presented at LREC2002, Las Palmas.

O'Madadhain, J., Fisher, D., Smyth, P. White, S., & Boey, Y. (2005). Analysis and visualization of network data using JUNG. *Journal of Statistical Software*. Retrieved July 7, 2008 from http://jung.sourceforge.net/doc/JUNG_journal.pdf

Pasin, M., Motta, E., & Zdrahal, Z. (2007). Capturing knowledge about philosophy. *Proceedings of K-CAP 2007 - The Fourth International Conference on Knowledge Capture*

Pease, A., Niles, I., & Li, J., (2002). The Suggested Upper Merged Ontology: A large ontology for the semantic web and its applications. *Working Notes of the AAAI-2002 Workshop on Ontologies and the Semantic Web*

Poli, R. (1995). Bimodality of formal ontology and mereology. *International Journal of Human-Computer Studies*, 43, 687-696

Rescher, N. (2007). The transformation of American philosophy. In *Is Philosophy Dispensable? And Other Philosophical Essays* (pp 99-114). (Frankfurt: Ontos Verlag)

Richardson, M., & Domingos, P. (2006). Markov logic networks. *Machine Learning* 62(1-2), 107–136

Sanger, L. (2008). Who says we know? On the new politics of knowledge. *The Edge*, 208.
Retrieved March 22, 2008 from http://www.edge.org/3rd_culture/sanger07/sanger07_index.html

Smith, B. (2003). Ontology. (In Luciano Floridi (ed.), *Blackwell Guide to the Philosophy of Computing and Information* (pp155-166). Oxford: Blackwell.)


Smith, B. (2006). Against idiosyncrasy in ontology development. *Proceedings of Formal Ontology and Information Systems 2006*

Stuckenschmidt, H. (2006). Toward multi-viewpoint reasoning with OWL ontologies. *Lecture Notes in Computer Science*, 4011

Trautwein, M. and Grenon, P. (2004). Roles: One dead armadillo on Wordnet's speedway to ontology. *Proceedings of the Second International WordNet Conference*, 341-346

CLASS BROWSER

For Project: inpho

Class Hierarchy 

owl:Thing

Idea

Idea_in_Non-Western_Philosophy

Idea_in_Arabic_and_Islamic_Philosophy

Idea_in_Chinese_Philosophy

Idea_in_Japanese_Philosophy

Idea_in_Western_Philosophy

Idea_in_Epistemology

Idea_in_Ethics

Idea_In_Applied_Ethics

Idea_in_History_of_Ethics

Idea_in_Metaethics

Idea_in_Metaphysics_of_Morals (13)

Idea_in_Moral_Epistemology_and_Moral_Reasoning (5)

Idea_in_Normative_Ethics (8)

Idea_in_History_of_Western_Philosophy

Idea_in_Logic

Idea_in_Metaphysics (2)

Idea_in_the_Specialized_Philosophies

Nationality (83)

Organization

Profession (905)

Publication (3)

Thinker (1613)

INSTANCE BROWSER

For Class: Idea_in_Metaphysics_of_Morals

Asserted

Inferred

:NAME

Constructivism

Error_Theory

Expressivism

Fictionalism

Moral_Antirealism

Moral_Naturalism

Moral_Non-Naturalism

Moral_Particularism

Moral_Realism

Neo-Kantianism

Quasi-Realism

Theological_Voluntarism

Value_Pluralism


Philosophy [collapse all] [expand all]

- ☒ aesthetics and philosophy of art
- ☐ african and african-american philosophy
- ☐ arabic and islamic philosophy
- ☐ chinese philosophy
- ☒ continental philosophy
- ☒ epistemology
- ☒ ethics
- ☒ feminist philosophy
- ☒ history of philosophy
- ☐ japanese philosophy
- ☐ judaic philosophy
- ☐ latin american and iberian philosophy
- ☒ logic
- ☒ metaphysics
- ☐ philosophy of action
- ☒ philosophy of language
- ☒ philosophy of law
- ☐ philosophy of mathematics
- ☒ philosophy of mind
 - ☒ artificial intelligence
 - ☒ consciousness
 - ☒ mental content
 - ☒ metaphysics of mind
 - ☐ emergence
 - ☐ free will
 - ☒ mind-body problem
 - ☐ personal identity
 - ☐ reduction
 - ☐ supervenience
 - ☒ perception
 - ☒ philosophy of psychology
- ☒ philosophy of religion
- ☒ philosophy of science and the sciences
- ☒ social and political philosophy

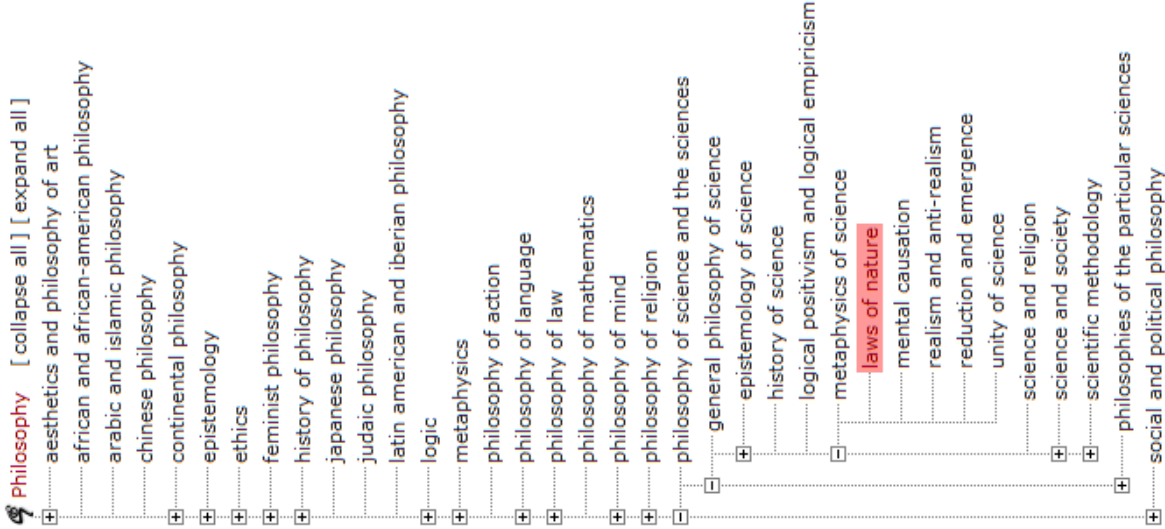
« myinpho (login)

- Philosophy [collapse all] [expand all]
- ⊕ aesthetics and philosophy of art
 - ⊖ african and african-american philosophy
 - ⊖ arabic and islamic philosophy
 - ⊖ chinese philosophy
 - ⊕ continental philosophy
 - ⊕ epistemology
 - ⊕ ethics
 - ⊕ feminist philosophy
 - ⊕ history of philosophy
 - ⊖ japanese philosophy
 - ⊖ judaic philosophy
 - ⊖ latin american and iberian philosophy
 - ⊕ logic
 - ⊕ metaphysics
 - ⊖ philosophy of action
 - ⊕ philosophy of language
 - ⊕ philosophy of law
 - ⊖ philosophy of mathematics
 - ⊖ philosophy of mind
 - ⊖ artificial intelligence
 - ⊖ consciousness
 - ⊖ mental content
 - ⊖ metaphysics of mind
 - ⊖ perception
 - ⊖ philosophy of psychology
 - ⊖ aspects of mind
 - ⊖ folk psychology
 - ⊖ **philosophy of cognitive science**
 - ⊕ philosophy of religion
 - ⊕ philosophy of science and the sciences
 - ⊕ social and political philosophy

This taxonomy of ideas is the outcome of the first iteration of our cycle of expert-provided structure, statistical analysis of the articles in the [Stanford Encyclopedia of Philosophy](#), and a small amount of human feedback. Iterations of this cycle will be used to develop the taxonomy. The approach is explained in more detail [here](#).

Click on any of the links on the left to explore the topics in the taxonomy. Click  to open SEP article on the adjacent term.

Instances:		Searches				
🔍	mental illness	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
🔍	bayesian epistemology	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
🔍	animal consciousness	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
🔍	self-reference	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
🔍	emergent properties	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
	performance	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
	indeterminacy of translation	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
	materialism	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
	semantic holism	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
	evolution of cognition	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
	rationality	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
Links:		Searches				
🔍	@qualia	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
	@perception	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
	@consciousness \cap psychology	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
	@theories of mental content	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
	@metaphysics of mind	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
🔍	@artificial intelligence	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
🔍	@evolutionary psychology	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
	@philosophy of science	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
	@language learning	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
🔍	@compositionality	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
	@logic \cap computation	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
	@philosophy of logic	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
🔍	@consciousness \cap intentionality	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
	@externalism	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
	@symbolic processing	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar
🔍	@animal cognition	🔍 SEP	noesis	🔍 8 SEP	Google	Scholar



You can navigate through the idea tree by clicking on the topic node links the left. Clicking on a node expands it into the available subtopics. If no terms appear below, on the right side of this page, then please follow links until some do.

For each term shown below on the left, please indicate its relationship to the topic node selected (laws of nature). You may skip any items you are unsure about. For more information about what you are being asked to do, please click [here](#).

- How should I decide the relatedness or relative generality of two ideas? (see 1-4)
- How hard should I think about idea pairs which seem odd?
- Can I review/revise my feedback?

Page 1 2 3 4 5 6 7 8 9 10 Add your own | Jump to submit button

unrelated	<input type="radio"/> <input type="radio"/> <input type="radio"/> <input checked="" type="radio"/> <input type="radio"/> <input type="radio"/> highly related	causal determinism	is more specific	laws of nature	<input type="button" value="reset"/>
<input type="checkbox"/> This is not a philosophical idea. ?					
<input type="button" value="SEP"/> <input type="button" value="noesis"/> <input type="button" value="8 SEP"/> <input type="button" value="Google"/> <input type="button" value="Scholar"/>					

unrelated	<input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> highly related	arguments for incompatibilism		laws of nature	<input type="button" value="reset"/>
<input type="checkbox"/> This is not a philosophical idea. ?					
<input type="button" value="SEP"/> <input type="button" value="noesis"/> <input type="button" value="8 SEP"/> <input type="button" value="Google"/> <input type="button" value="Scholar"/>					

unrelated	<input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> highly related	compatibilism		laws of nature	<input type="button" value="reset"/>
<input type="checkbox"/> This is not a philosophical idea. ?					
<input type="button" value="SEP"/> <input type="button" value="noesis"/> <input type="button" value="8 SEP"/> <input type="button" value="Google"/> <input type="button" value="Scholar"/>					

unrelated	<input type="radio"/> <input type="radio"/> <input type="radio"/> <input checked="" type="radio"/> <input type="radio"/> <input type="radio"/> highly related	descartes' physics		laws of nature	<input type="button" value="reset"/>
<input type="checkbox"/> This is not a philosophical idea. ?					
<input type="button" value="SEP"/> <input type="button" value="noesis"/> <input type="button" value="8 SEP"/> <input type="button" value="Google"/> <input type="button" value="Scholar"/>					

unrelated	<input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> <input type="radio"/> highly related	miracles		laws of nature	<input type="button" value="reset"/>
<input type="checkbox"/> This is not a philosophical idea. ?					
<input type="button" value="SEP"/> <input type="button" value="noesis"/> <input type="button" value="8 SEP"/> <input type="button" value="Google"/> <input type="button" value="Scholar"/>					

is more specific than
is more general than
is as general as
incomparable/either

